CONFERENCE PROCEEDINGS

RAND

Research and Development Initiatives Focused on Preventing, Detecting, and Responding to Insider Misuse of Critical Defense Information Systems

Robert H. Anderson

National Security Research Division

Research and Development Initiatives Focused on Preventing, Detecting, and Responding to Insider Misuse of Critical Defense Information Systems

Results of a Three-Day Workshop

Executive Summary

On August 16-18, 1999, approximately 40 researchers and government research sponsors involved with information system security met at RAND, Santa Monica CA, to address and recommend technical research and development initiatives focused on mitigating the insider threat. The workshop was sponsored by NSA/R2, DARPA/ISO, and the Army Research Laboratory.

Although the workshop's main purpose was to propose technical research initiatives, it was clear to all participants that enabling policies are required in order for the results of insider threat research to be effective. Policies and procedures needed to form an environment for mitigating the insider threat include: *guidance and requirements for researchers from the legal and law enforcement communities* regarding the attribution, collection, maintenance, processing, and storage of data in a manner that allows proper forensic analysis, and a trail of custody to permit later legal prosecution; *clear definitions regarding what constitutes "critical assets"* on a system to be protected against insider misuse; *clarity about the definition of an "insider;" cost/benefit analysis* to help determine whether the cost to personnel and organizations, as well as dollar cost, of new security procedures are worth the security benefits obtained; *plans for technology transfer*; and *support of multiple, diverse, concurrent approaches.*

Another theme pervading the workshop was that the insider threat is not unique unto itself; it is part of overall information system security. Many of the research initiatives proposed cross-cut with "traditional" information security measures. The insider threat particularly raises issues of personnel screening/testing procedures and other policies that were not the focus for this workshop. It was clear, however, that there are specific insider threat

problems that can be addressed by information system initiatives and technologies. To the extent possible, participants tried to emphasize the insider-unique characteristics of the approaches that were discussed.

Workshop discussions and breakout groups used the following categories in describing possible research initiatives: those addressing threats and vulnerabilities, prevention, detection, and response. We summarize here the key recommendations in each of those categories.

Threats and vulnerabilities. A threat has as elements motivation, opportunity, vulnerability in the target system, and skill adequate to exploit the vulnerabilities. The primary research initiative aimed at countering threats and vulnerabilities was:

• *Develop an insider trust model.* The category "insider" includes, among many others, part-time employees, temporary employees, and occasional insiders. What degree of trust should be given to each such category? How can this level of trust be represented explicitly – perhaps even in machine-readable form – so that various processes within the system can take cognizance of those varying levels of trust?

Prevention. The specific research initiatives discussed by the prevention group centered on the goal of creating system security through a set of components, no one of which is sufficient or completely trusted, and making that set of components usable in practice. It was felt that overcoming obstacles to use of available components is itself a research topic. Components out of which prevention can be constructed would result from these proposed initiatives:

• *Create a set of authentication components.* Such components must apply to a multi-tier transactional environment. They must bind keys and tokens to users. They must meet performance and operational standards such as exceeding 1000 transactions per second. They must include practical revocation and recovery. And they must integrate with the various applications in use within the system.

• *Develop access control components, for both normal and privileged users.* We need finer-grained access control to reduce vulnerability to the insider threat, but such granularity is expensive. Most existing controls are for single platforms only. We need inter-platform access-control management. A key research need is means of reducing the management cost of implementing and maintaining access controls. Privileged users are a special case: Whoever controls the access controls is the *primary* insider issue. The question is how to mitigate risks from malicious system or security administrators, and other privileged users. One aspect of access control is *malicious code containment*, so that access to important system components is limited for malicious code introduced to a system.

• *Develop trusted paths to security systems.* Even if system integrity is assured, there are alternative means by which a malicious user can gain system privileges – for example, by spoofing the graphic user interface by which a user is asked for his login and password. A potential solution is obtaining an assured trusted path to the security system. The availability of such a path is a fundamental component of a secure architecture.

Detection. The following recommendations deal with the problem of detecting an insider's misuse:

• *Develop profiling as a technique.* A "user profile" contains information that characterizes users or agents, and may help to distinguish one user from another. (The term "user" here may refer to a process, command, function, etc.) The profile may contain such information as those files and processes normally accessed; periods of time that user is normally logged in; keystroke patterns; and a wide variety of other attributes. With such profiles available, it is then possible to develop *anomaly detection* for unauthorized results that have not been previously encountered.

• *Detect misuse of applications.* Persons concerned with the "insider problem" often focus on insiders' inappropriate access to data. However, another useful indicator of misuse is inappropriate use of processes and applications within the system. (A most obvious example is an attempt to use system processes to obtain root access.)

• *Provide traceability for system-object usage.* It would be highly desirable to develop audit trails of the path of usage of system objects such as files and programs. If misuse is suspected, these audit trails could be investigated to determine the set of users accessing an object, or passing a system object among them.

• *Identify critical information automatically.* Contemporary information systems contain a constantly shifting collection of gigabytes of data and information. It seems unlikely that static lists of critical files and processes will remain relevant. This research recommendation addresses the question: Can procedures be developed by which critical information within a system is identified automatically?

• *Design systems for detectability.* Detecting misuse by an insider is very difficult. It could be made easier if the inherent architecture and design of an information system was created with detectability of misuse in mind. This research task aims at creating system design principles that have detectability as one of their primary features. One important aspect of detectability is determining unauthorized changes due to physical access. A distinguishing characteristic of insiders is that they often have physical access (at varying levels) to system facilities. As such, they can make physical changes to the system -- such as installing temporarily another client or host computer on the network, rerouting a cable, or plugging in a modem -- that are precluded to outsiders. It is therefore important to monitor for and detect such physical changes to a system, preferably in real-time.

Response. Research initiatives to aid in responding to an insider misuse event include:

• *Incorporate practical autonomic system response into production systems.* Insiders can carry out destructive acts in seconds. Spontaneous system responses to detected incidents may at times be required in order to thwart or mitigate damage in similar time scales – which would imply use of automated response procedures. The question is: How can these be developed so that they can be effective while not causing more harm than good. In particular, if a malevolent insider knows such an autonomic system is in use, it could be used to harm the system itself. If system developers built capabilities relevant to response into system architectures, it would be much easier to initiate effective response, perhaps autonomically, within critical information systems.

• *Develop data correlation tools, including data reduction for forensics, and visualization tools focused on internal misuse.* In responding to a potential or actual misuse incident, data from multiple sources must be correlated and analyzed very quickly, data captured and stored for forensic analysis, and complex patterns visualized and analyzed to obtain understanding so that response can proceed. Such tools are currently lacking, especially when confronted with an insider threat – in which case many traditional sources, such as those associated with a firewall, may be unavailable or irrelevant.

• *Develop a capability for surveillance of non-networked components.* Not all information "system" components are on-line and capable of being queried for relevant data. Telephone logs, timecard data, various hardware settings, and so on may be relevant to determining a response, but be unavailable to correlation and fusion activities. To enable rapid and relevant response, it is highly desirable that all relevant data be collectable in real-time.

• *Consider deception technologies specifically applicable to the insider threat.* Deception has long played a vital role in military operations -- both offensive and defensive. It has tantalizing capabilities to aid in determining aspects of insider misuse that are otherwise hard to discover – in particular, the interests and intentions of a malicious insider, and that person's level of sophistication in understanding and using system facilities. Research should be conducted on the appropriate uses of deception to aid in understanding attributes of the malicious insider, and to divert him or her from critical system files and processes.

Participants agreed that this workshop should be the first in a series. A potential topic for a second one is methods for correlation and fusion of information received from various sensing and detection mechanisms. Such correlation is needed so that insider (and other) misuse can be discovered quickly so that appropriate rapid response can be taken. It is also vital that there be cross-fertilization between researchers concentrating on mitigating the insider threat, and those working on other aspects of information system security. The problems are interrelated, so a number of common research and development approaches should be pursued.

Background

It has become widely acknowledged that "insider misuse" is one of the greatest threats to -- and obstacles in achieving -- "information assurance" within critical Defense information systems.[1] In recognition of this, an Integrated Process Team (IPT) was formed recently within OASD/C3I, under the leadership of Mr. Tom Bozek, to create a DoD Insider Threat Mitigation Plan. The final report from the IPT has been published.[2] We assume the reader of this current document is familiar with the definitions and findings of the IPT report. Recommendation 1.2 (see Appendix B, pg. B-2) of that report states, "Conduct recurring workshops on technological approaches to mitigating the insider threat and reducing information system vulnerabilities" – with action for implementation assigned to C3I and NSA. The workshop reported on here can be considered the first step in implementing that recommendation.

On August 16-18, 1999, a three-day invited workshop (expected to be the first in a series) was held at the facilities of the RAND Corporation, Santa Monica CA, specifically to address and recommend technical research and development initiatives focused on mitigating the insider threat. The workshop was sponsored by NSA/R2 (Dick Brackney), DARPA/ISO (Sami Saydjari), and the Army Research Laboratory (Paul Walczak). A complete list of attendees is shown in Attachment 2. The agenda for the workshop is given in Attachment 3.

The IPT report's recommendations focus on near-term steps that can be taken, preferably at modest time and cost, to mitigate the insider problem. This workshop can be considered a supplement to that report, focusing instead on more significant, specific R&D initiatives that can be taken to address the insider problem at a more fundamental level.

This report summarizes the workshop's key findings and recommendations. Four categories of breakout sessions were used during the workshop: on the threat, prevention, detection, and response. We use those same categories here in describing the workshop's results.

---

[1] The insider problem is pervasive in all information systems. We concentrate here on Defense systems, although much of this discussion and set of recommendations applies more widely to other government systems, critical infrastructure systems, and private commercial systems in general.

[2] *DoD Insider Threat Mitigation Plan: Final Report of the Insider Threat Integrated Process Team*, June 1999 FOUO.

Prologue: Policy and Precursors

The workshop stressed tangible research and development initiatives that could significantly mitigate the insider threat. However, participants felt some important policy and procedure initiatives were required to create an environment in which such R&D could be effective. These precursors to R&D are summarized below.

PR1: The research community needs guidance and requirements from the legal and law enforcement community regarding the attribution, collection, maintenance, processing and storage of data. In the IPT report, there are frequent references to developing an untamperable audit trail, so that detected misuse by insiders can stand up to legal scrutiny in court. Before researchers can develop technical means of collecting and storing data indicating insider misuse, they need clear guidance from those who perform forensic analysis and later legal prosecution regarding the requirements to be met by such data, and by the trail of custody within which it is kept.

PR2: Clear definitions are needed regarding what constitutes "critical assets" on a system to be protected against insider misuse. It was clear to participants that it is difficult if not impossible to protect all data on a system against all forms of misuse. Therefore, each information system should be accompanied by a clear statement – available to users and system administrators alike – regarding what assets on the system are regarded as critical, and whose misuse will be treated as a punishable offense. Perhaps there might be several such categories or gradations. But without such clear guidance, it is difficult to focus preventative and detection measures on those assets that are deemed critical, and without such focus the efforts at data collection and analysis are likely to be overwhelmed by the "noise" of lesser offenses, mistakes, and the like. Once such guidance regarding what is critical information is available – preferably in machine-processable form – it might be possible to institute automatic or semi-automatic means of assessing the criticality of information as it is created, manipulated, and stored within an information system. (See recommendation **D4**, below.)

As a corollary to this requirement, it is necessary for DoD to provide explicit guidance regarding which information systems as a whole are deemed critical – again, so that limited R&D and operational resources can be focused on those whose operation and data are most valuable and that need to be most protected.[3]

PR3: Be clear about the definition of an "insider."

One difficulty in discussing the insider threat was lack of a common understanding and definition of the term "insider threat." It was clear from discussions in the workshop that the term was like a chameleon – its "color" can change depending upon both the insider and the insider's environment. It is likely that "insider threat" is too broad of a term to usefully consider specific problems and their potential solutions. However, a number of possible insider threat descriptors or metrics were captured at the workshop that might help narrow future discussions.

Defining the environment:

Physical access – "spatial perimeter" Does the particular problem include physical access to critical system components? Depending on definition of the problem, physical access may extend beyond control of information systems. For example, can the insider turn off the air conditioning or power?

Cyber access – "logical perimeter" Does the particular problem include some type of cyber defense perimeter such as routers, firewalls, and protective applications?

Technical environment The technical details of the environment will likely have an effect on the insider threats that might present themselves and the way in which an insider operates.

Law enforcement environment Are there any restrictions or requirements about preparation for or response to an insider event that are important for law enforcement activity? An example is integrity of evidence and chain of custody. A particular concern is the conflicts this item may have with the technical environment.

Defining the insider:

Normal – Abnormal – Malevolent What is the intent of the insider? "Normal" insider activity will likely not pose a threat. "Abnormal" insider activity may include routine errors that could

[3] This recommendation parallels the discussion in section 2.5.1 of the IPT report: "Declare What is Important." Participants felt this point merited special attention.

cause a weak system to have problems or private information to be unintentionally revealed. "Malevolent" insider activity includes that with malicious intent.

Novice – Sophisticated What level of skill does the insider possess? This includes factors such as the amount of preparation taken, the adversity to detection, etc.

Knowledge of internal environment How much does the insider know about the workings of the system inside the spatial or logical perimeter? Various degrees of insider knowledge and understanding were discussed at the workshop.

Innate privileges What privileges does the insider have both physically and administratively? This is related in part to knowledge of the internal environment.

Degree of Affiliation. The relationships among insiders, their specific roles/privileges and the nature of their affiliation with the system's organization helps to provide greater understanding of the threat. Insider threat characterizations that differentiate permanent, part-time, temporary, outsource, former insiders, and system developers, help to identify unique opportunities, motivations, skills and vulnerabilities that might otherwise be overlooked.

Human – Code/Hardware Is the insider a person, software, firmware, or hardware? There was much discussion at the workshop that an insider did not have to be a person. For example, malicious code was viewed to be important and currently a poorly addressed insider threat.

The IPT report defines an insider as anyone who is or has been authorized access to a DoD information system whether a military member, a DoD civilian employee, or employee of another Federal agency or the private sector. It provides a table (in section 2.2) that includes within this definition contractors, network-connected users at colleges and universities, foreign partners, and vendors and suppliers. Although our workshop used the IPT report as source material, considerable discussion focused on other definitions, such as persons having a "trust" relationship with the organization in question. If insider definitions are so broad as to encompass 100,000 employees of Microsoft, Sun Microsystems, or Cisco (as vendors to DoD), those definitions lose most of their force and focus. If limited R&D resources are to be focused explicitly on the "insider," what definition should be used – at least by the research community – to focus their efforts on truly critical breaches of trust by "insiders"?

Among the statements made at the workshop regarding the definition of "insider" were these:

• "Insider" can be characterized based on technology management roles (and the trust relationship with which each role is accorded), as well as process accessibility inside the organizational domain of control (with respect to the degree of confidentiality, integrity, and availability accorded). For example, an insider who attempts to achieve an objective by violating confidentiality requirements is more likely to engage in covert activities, hoping to avoid detection by the enforcement/oversight mechanisms within the victim organization. By contrast, an insider whose objective is to interfere with corporate business processes may execute denial-of-service actions expecting that his/her acts will become readily apparent, but not attributable.

• Three aspects possibly differentiating an insider from an outsider are: (1) knowledge of the internal environment; (2) speed of attack, based on availability; (3) relative ease of accessibility.

• What constitutes an "insider" differs between law enforcement and technical personnel. For law enforcement, the Insider Threat involves use of internal knowledge by people authorized to use the system; the threat involves a violation of trust within the group of persons given that trust. For technologists, the Insider Threat is that which happens inside the security perimeter, i.e., behind the firewall, from the perspective of the operating system components. The "insider" is anyone influencing a process operating inside a group of trusted hosts or within system protection domains.

• Military personnel can be especially prone to negative motivations that lead to their becoming an insider threat, due to lifestyle restrictions imposed as part of their "service contract;" 7x24 service; attestation to a new set of "military values;" the sacrifice of some liberties; pressures toward conformity – social, moral, cultural; arduous work and living conditions; fatigue; boredom. Additionally, aspects of military culture that denigrate support and technical roles, especially within combat organizations, can exacerbate vulnerabilities to the insider threat. These cultural conditions can stimulate malicious insider activity by service members as well as the exodus from active duty of knowledgeable military personnel who may now bear hostilities as "former insiders." Several of these factors can lead to disillusion.

• The military's Total Force concept, using reserves and the National Guard, creates "part-time" insiders. Other "insiders" are created by Combined Warfare, comprising alliances, coalitions, and host-nation support .

• There is accelerated information technology infusion underway (e.g., the First Digitized Division; Army XXI; Navy IT 21). In addition, a torrent of Y2K "fixes" will introduce new bugs into an immature, unproven information technology environment. Software agent/mobile code technology may represent a non-human variant of the "insider threat," and outsourcing beyond existing manning to sustain operations creates "partnering insiders."

For all these reasons, insider misuse at DoD must focus on relevant implications across a spectrum of military operations, functional capabilities, and technology development programs, to include: strategic to tactical; peacetime to high-intensity conflict; sustaining base to battlespace; contractor to flag officer to SECDEF; sensors to mobile wireless communications platforms to computer-based information systems.

PR4: Cost/benefit analysis is vital. Implementation of many of the recommendations in the IPT report, and results likely to emerge from new R&D initiatives, almost invariably place additional constraints on users or systems. Such constraints may well negatively impact productivity. They may, in fact, "turn off" users and sysadmins to the point where they develop alternatives to the use of these protective/detective/reactive measures. A serious cost/benefit analysis must be done for the recommendations of the IPT report, weighing potential safety/security benefits against personal and organizational impacts on productivity and effectiveness. Similarly, this same analysis must be done before implementing outcomes from the R&D initiatives proposed in this document. We realize, however, that cost/benefit analysis is very difficult when the "benefit" of security is somewhat intangible, as is the "cost" to personnel and organizations (not to mention dollar cost) which is equally difficult to measure. Perhaps economists, organizational researchers and researchers in other such disciplines should be marshaled to address this difficult analysis problem explicitly.

PR5: Detection should include (but not be limited to) host-based information. Many information assurance measures have been tailored to preventing outsiders from obtaining access to systems. Such measures have therefore concentrated on "firewalls" placed on entrances

to internal networks, or on network monitoring tools that examine network packets for abnormal patterns or for "signatures" of known attack techniques. However, the insider is not coming in through external portals; therefore, it is important to consider the "host" or "client" computer (e.g., desktop computer or workstation) used by an insider in monitoring and detection systems aimed at detecting insider misuse.

PR6: Facilitate and plan for technology transfer. It is a commonplace in the research community to lament that most of what is known about effective security practices in the design, implementation, and operation of information systems is not, in fact, being used in contemporary operating systems, application programs, and networks. In fact, much of the output of the computer security research community over the past several decades has not found its way into COTS products heavily used by DoD and others. There appears to be a gap between "research and early prototype development" and the next step – serious development for widespread test and evaluation. There are a number of approaches being investigated toward bridging this gap, such as DARPA's Technology Integration Center (TIC), but more needs to be done on this issue.

PR7: There are effective approaches to the insider threat that are not technology-based. It was noted by workshop participants that – although this workshop concentrated on technology R&D approaches – there are other means of combating the insider threat. (In particular, approaches involving training, education, responsible system administration, management initiatives, and so on, are mentioned in the IPT report.) An example stressed by one participant was greater use of polygraph tests for insiders on critical information systems. Such tests have been used to uncover patterns of misuse that would otherwise go undetected.

PR8: Develop a database of case studies of insider misuse. To guide research and other policy/procedural initiatives for mitigating insider misuse, we need better data on the aims, means, level of sophistication, and degree of success of insiders that have misused important DoD information systems. Only by analysis of such data can we be assured that initiatives to stem misuse are aimed in appropriate areas, and that important misuse methods have not been overlooked.

PR9: There is no silver bullet. We need multiple, diverse, concurrent approaches. It was clear to participants that "defense in depth" is vital. Multiple approaches must be used,

and to the extent possible the findings and sensings from multiple probes and detectors, and multiple protection systems, must be coordinated and correlated to understand when a vulnerability exists and is being exploited. Especially in the case of insiders, there need be no explicit "vulnerability" at all; an insider may have valid access to critical, sensitive data, and in one download or e-mail attachment create a very significant security breach. In that sense, no one should expect the insider problem to be "solved" – at best, we can protect, detect and react in such a manner as to mitigate the problem to the extent possible.

In the explicit R&D recommendations that follow, we attempt where possible to state specific research objectives, success metrics, and open research problems. We also try to characterize factors that distinguish the unique insider-related problems and approaches from those of more general information security. However, there is considerable overlap between security approaches needed for general information security and those specifically tailored toward to the insider threat. It is often the case that general security awareness, training, education, protection, detection, and reaction measures will assist directly with the insider problem also.

Insider Threats and Vulnerabilities

The workshop breakout session on insider threats and vulnerabilities modified a JTF-CND (Joint Task Force – Computer Network Defense) vugraph to provide an overview of the distinctions among an incident, an attack, and a specific event (Figure 1). This figure also indicates the roles of an attacker's motivation combined with access (providing an opportunity), plus use of skills and tools, and shows the role of detection technology when an event has occurred. As indicated in the figure, a threat comprises four elements:

- Motivation
- Opportunity to pursue motives
- Vulnerability in the target information system, and
- Skills adequate to exploit the vulnerability.

Any threat analysis is under at least one constraint: the "impact threshold" – that is, determining when risk exposure is too costly to ignore.

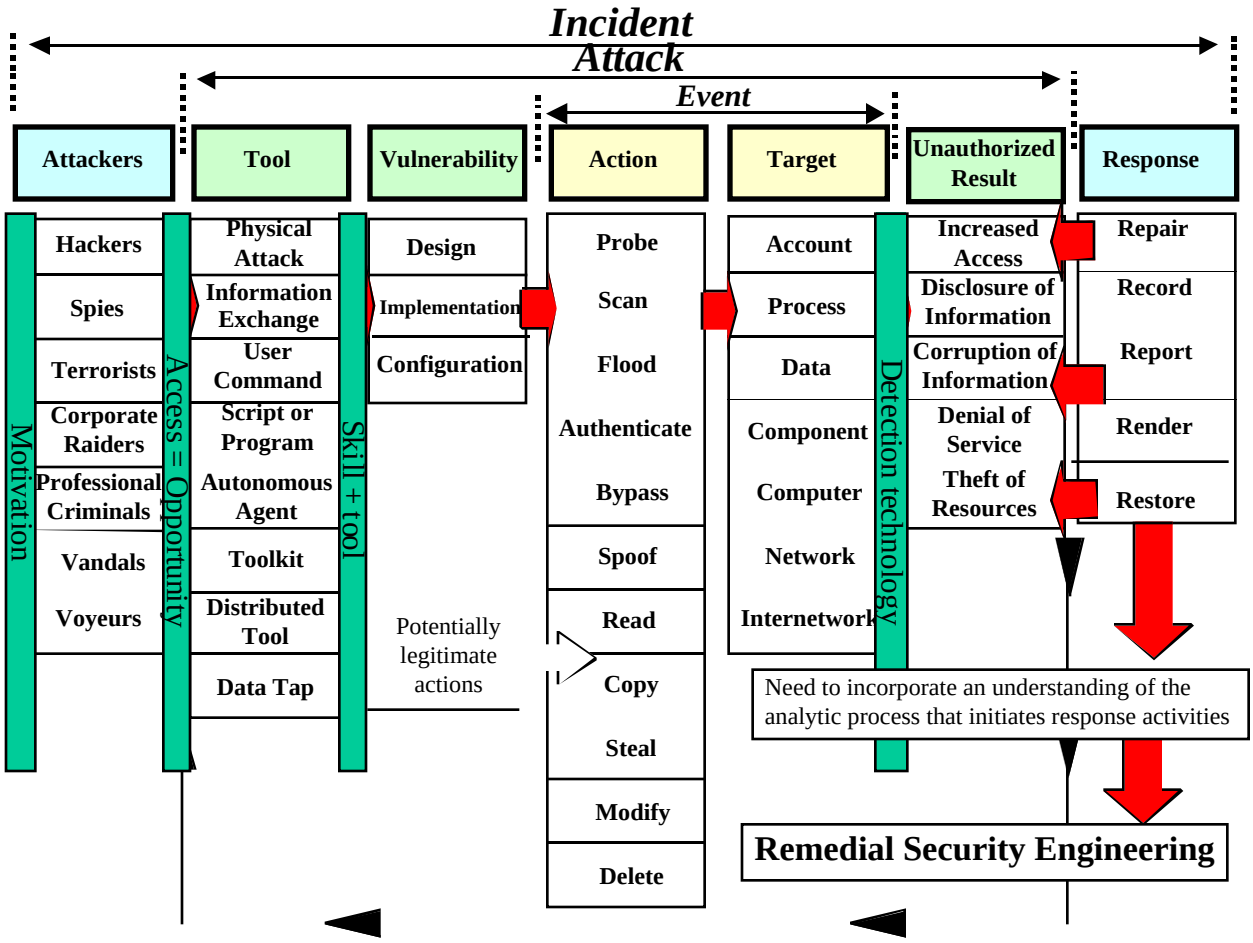

Fig. 1 -- Characterizing an Information System Security Incident
Source: Significantly modified from original
Joint Task Force – Computer Network Defense document

The "threat" breakout group characterized the impact of an insider threat in three categories:

- The degree of "reach" (i.e., exposure, or risk), to critical components and to sensitive data, and the associated cost of restoration
- The frequency and intensity of occurrence, and the cost of detection
- Business and technical impacts of the threat.

Insider threat may be associated with at least each of the following roles, with varying roles implying varying degrees of risk:

- Authorized user
- Network administrator
- Systems administrator
- Information security officer
- CERT personnel
- Systems maintenance personnel
- Building maintenance
- Building security.

The threat group highlighted four areas in which research activities should be conducted:

T1: Develop reactive configuration controls, in which an unauthorized result is mapped back to a specific type of threat. When the results of a misuse event is detected, there should be means for deciding the nature and severity of the threat represented by that event. Depending on the nature and severity of the threat, it may be possible to automatically reconfigure the system to react in such a manner that further threat is minimized.

Research objectives:

(1) Characterize the insider threat.

Success metrics:

(1) Ability to distinguish a variety of threat conditions over a range of environments.
(2) Demonstrate ability to respond to tune technology to achieve high-fidelity sensitivity to anomalous insider activity.

Insider distinguishing factor:

Certain routine insider activity might be interpreted as malicious behavior using the "outsider" perspective for information assurance

Open research problems:

(1) Identify insider misuse characteristics with respect to various operating environments, topologies, applications and roles.
(2) Compare and contrast insider vs. outsider ability to achieve adverse, unauthorized results.
(3) Demonstrate that computer security events can be traced back to specific insiders, even when the insider attacks from "outside" the trusted domain.

T2: Develop an insider trust model. The category “insider” includes, among many others, part-time employees, temporary employees, and occasional insiders. What degree of trust should be given to each such category? How can this level of trust be represented explicitly – perhaps even in machine-readable form – so that various processes within the system can take cognizance of those varying levels of trust?

Research objectives:

(1) Develop a model of trust that envelops the full breadth of roles for which an organization authorizes degrees of technical configuration control privilege.

Success metrics:

(1) User/process classes, roles and use-cases are documented in some hierarchical, form, either as a data schema or a matrix of conditions, that clearly delineates the privileges authorized to a user for any system resource during any potential system state.

Insider distinguishing factor:

The degree of certain attributes of the trust relationship is the key distinguishing attribute separating a notion of insider from outsider.

Open research problems:

(1) Develop a characterization schema describing representative user (insiders) roles and privileges in a DoD environment, that covers the full spectrum of military operations.
(2) Differentiate between the various roles identified in (1) and develop parametric sensitivity criteria that can be used to recognize attempted unauthorized escalation of privilege within the timeframe between vulnerability exploitation and achievement of adverse objective (potential security breaching event).

T3: Develop means to map users to unauthorized results. If a system anomaly is detected, did an insider cause it? If so, which insider? Can we develop means by which outcomes can be traced back to perpetrators, especially when they are insiders? (Note: This recommendation is closely related to detection recommendation **D3**, below. See that recommendation for specific details on research objectives, success metrics, and open research problems.)

T4: Identify signatures of unauthorized results. What, in fact, are the indicators that a system has been misused? We cannot know the level of threat we face if we cannot identify unauthorized results when they are present. Developing "signatures" of unauthorized results (e.g., based on case studies and analysis of misuse possibilities) is one research approach.

Research objectives:

(1) Focus insider misuse detection upon unique vulnerabilities presented by the insider threat (lead into detection recommendations).
(2) Develop an understanding of insider patterns that can be detected by machines.

Success metrics:
(1) Identification of patterns of misuse unique to insider threat.

Insider distinguishing factor:
Self-explanatory; objective here is to find insider-distinguishing patterns of misuse.

Open research problems:
(1) Prove that sensors can reliably alert to specific examples of signatures identified above (related to detection recommendations elsewhere in this report)

Prevention

The prevention breakout group chose to identify components that are necessary for solving the insider threat: authentication, access control, system integrity, trusted path, and attribution. These components were each expanded upon and provided below as P1 through P5.

Discussions in this group identified no key "insider distinguishing factors" that were unique for the prevention view of the insider threat (especially in the component view provided in this section), or that were particularly different from other (vulnerabilities, detection, response) workshop discussion. Therefore, this characteristic is not identified individually for P1 through P5. However, several specific points were captured on this topic:

- Only the insider can abuse authority or trust (by definition), and knowledge is not equivalent to trust in this context
- Insiders have superior knowledge of asset value, thus increased vigilance is required on high-value assets
- One can arrest an insider, while arresting an outsider can be problematic. (In a prevention context, possibility of an "arrest" may deter inappropriate insider activity. Also, it can complicate the related preventative measures that are necessary, such as providing sufficient and verifiable evidence.)

Mechanisms for preventing insider abuse are not unique to the insider threat; however, these traditional mechanisms must be amplified and enhanced in order to address this threat.

P1: Develop authentication components. The ability to uniquely identify entities in a system and between systems with a high level of assurance provides a foundation for one class of solutions to the insider threat. If entities are identified with certainty, their activity can be

tracked, analyzed, and controlled. In addition, without strong authentication systems, all methods of detection and response are more likely to be subverted and rendered ineffective.

Some authentication technologies are beginning to become available with a limited scope and are not intended to scale to complex systems or systems with complex interactions. The problem is further complicated in that a sophisticated insider may have the ability to subvert authentication components that might otherwise flag inappropriate activity.

Research objectives:

(1) Extend technologies to work in multi-tier transactional environments.
(2) Ability to bind keys and tokens to users.
(3) Strong authentication that can scale for increasing transaction rates.
(4) Ability to include practical revocation and recovery.

Success metrics:

(1) Systems that meet performance and operational standards as necessary for users, such as ability to exceed 1000 transactions per second.
(2) Number of strong authentication tools widely available in COTS.

Open research problems:

(1) See research objectives.

P2: Develop access control components. The ability to provide access control is fundamental to controlling the insider threat; however, existing techniques have focused on limiting access at a very high level such as system passwords, user passwords, and for single platforms only. To be properly addressed, insider access must be controlled with fine detail such as per file, per transaction, per packet, etc. Procedures must take advantage of such access control to properly prevent and minimize insider problems. Access control must also operate between different platforms.

Research objectives:

(1) Development of finer-grained access control that is affordable.
(2) Inter-platform access control management.
(3) Reducing management cost of implementation and maintenance of access controls.
(4) Developing new types of access control that reduce vulnerability to trusted insiders.

Success metrics:

(1) Granularity of access control.
(2) Number (diversity) of access control methods.

Open research problems:

(1) See objectives.
(2) Expert system based access control automation that is able to translate natural language policy statements into machine-level policy.
(3) Meta-access control system for cross-platform access management.
(4) Robust and secure mechanisms for providing access control in order to prevent it becoming a single point of failure.
(5) Ability to prevent insider misuse from security administrators and other privileged users.

P3: Develop system integrity components. There are numerous hacker "tool kits" and other types of malevolent software available today that violate system integrity: for example, Back Orifice 2000 (BO2K), "Root Kit", and the Melissa virus. It is possible that there are "Trojan horses" being inserted into system code as part of Y2K fixes. Additional tools for system integrity checking and verification must be developed to address these threats.

Research objectives:

(1) Malicious code detection
(2) Arbitrary corruption prevention
(3) Develop boot sequence integrity
(4) Total system configuration management (for both hardware and software).

Success metrics:

(1) Number of COTS systems that implement thorough integrity checking.

Open research problems:

(1) See research objectives.

P4: Develop a bidirectional trusted path to the security system. Even if system integrity and authentication are assured (see recommendations **P1** and **P3**, above), there are alternative means by which a malicious user can gain system privileges – for example by spoofing the graphic user interface by which a user is asked for his login and password. A potential solution is obtaining an assured trusted path to the security system. The availability of such a path is a fundamental component of a secure architecture and involves authentication in both directions of a transaction. Providing this path requires both a systems viewpoint and attention to component operation detail.

Research objectives:

(1) Develop cross-platform trusted paths, both ways
(2) Develop two-way trusted paths in distributed systems
(3) Find ways to make trusted path concepts and techniques widely available in security architectures
(4) Find ways to instruct current-generation system developers in security techniques that are known, but not being applied.

Success metrics:

(1) Number of COTS systems implementing thorough trusted path procedures.

Open research problems:

(1) See research objectives.

P5: Develop attribution components. We need to be able to attribute specific actions to individual users. Approaches to this problem include watermarking (placing subtle or invisible marks in objects that attribute those objects' origin), and fingerprinting (watermarking that identifies where the object has been). These techniques are described in more detail in recommendation **D3**, below, and will not be described further here. We note, however, that the specific insider-related research problem is that an insider may have access to the attribution mechanisms, so they must be hardened against such insider misuse. A potential solution strategy is the use of strong cryptography within the attribution mechanism.

Detection

The following workshop recommendations deal with the problem of detecting an insider's misuse.

D1: Develop profiling as a technique. A "user profile" contains information that characterizes users or agents, and may help to distinguish one user from another. (Note: the term "user" here may refer to a process, command, function, etc.) The profile may contain such information as those files and processes normally accessed; periods of time that user is normally logged in; keystroke patterns; and a wide variety of other attributes. With such profiles available, it is then possible to develop *anomaly detection* for unauthorized results that have not been

previous encountered. Signature-based approaches only catch events that are previously recognized as malicious.

Research objectives:

(1) To discriminate between normal and anomalous behavior for a given user
(2) To be able to discriminate among users.
(3) To create technology that can identify new insider-initiated misuse.

Success metrics:

(1) Minimize number of false positives compared with correct detections per experiment
(2) Minimize performance overhead for obtaining and using user profile data.
(3) Using red-teaming experiments to inject misuse signatures not accessible (not in assurance system dataset) by targeted victim, demonstrate ability to recognize anomalous insider behavior.

Insider distinguishing factor:

Ability to collect user profile data is unique to the insider problem.

Open research problems:

(1) What are the best (sensor) sources of data?
(2) Feature extraction problems
(3) Best algorithms for detection
(4) Fusion/correlation of diverse information collected
(5) Scientific evaluation and comparison of techniques
(6) Design of contrastive experiments.

D2: Detect misuse of applications. Persons concerned with the "insider problem" often focus on insiders' inappropriate access to data. However, another useful indicator of misuse is inappropriate use of processes and applications within the system. (A most obvious example is an attempt to use system processes to obtain root access.)

Research objectives:

(1) Detect insider misuse of given resources and privileges
(2) Develop application-level sensors and detectors for misuse
(3) Go beyond access controls in user monitoring
(4) Generalize profiles to applications.

Success metrics:

(1) Minimize number of false positives compared with correct detections per experiment
(2) Minimize performance overhead for obtaining and using application access data.

Insider distinguishing factor:

This is a higher layer of detection that is specifically applicable to insiders, since system applications and processes are widely available to them.

Open research problems:

(1) Develop techniques for program profiling: What constitutes valid or invalid access to a program?
(2) Attempt to apply this detection technique within commercial operating systems
(3) Develop application-specific misuse detection
(4) Examine cases of insider misuse and develop a weighted threat model or matrix; what pattern of program accesses are an indicator of potential misuse?
(5) Develop auditability of object accesses; i.e., how to develop a tamperproof audit trail of access to such system objects as programs and processes.

D3: Provide traceability for system-object usage. It would be highly desirable to develop audit trails of the path of usage of system objects such as files and programs. If misuse is suspected, these audit trails could be investigated to determine the set of users accessing an object, or passing a system object among them.

Research objectives:

(1) Be able to determine who uses what, when, and how
(2) Detect suspicious exfiltration of data, programs, and intellectual property
(3) Provide object-centric traceability (e.g., by embedding an audit trail into documents and objects themselves).

Success metrics:

(1) Develop the capability to trace non-graphic objects (ability to trace graphic and audiovisual objects is being developed for intellectual property protection purposes by the entertainment industry)
(2) Minimize performance overhead for using such trace techniques.

Insider distinguishing factor:

This is quite specific to the insider problem, since the vast majority of uses of inside system resources is by insiders.

Open research problems:

(1) Mandatory watermarking of objects
(2) Embedding audit trails in objects
(3) The capability of applying these techniques to all kinds of system objects, such as text, graphics, source and binary code
(4) Retrofitting COTS software to enable the watermarking of intellectual property
(5) Developing appropriate algorithms to implement these techniques

(6) Developing appropriate infrastructure, such as RCS, compilers, and application wrapping.

D4: Identify critical information automatically. As mentioned earlier (see prologue item **PR2**, above), it is important to focus R&D and operational measures on those items within an information system that are critical. Contemporary information systems contain a constantly shifting collection of gigabytes of data and information. It seems unlikely that static lists of critical files and processes will remain relevant. This research recommendation addresses the question: Can procedures be developed by which critical information within a system is identified automatically?

Research objectives:

(1) Machine recognition of critical, possibly classified information by its content
(2) Development of machine processible classification guides (to be used by the automated recognition procedures).

Success metrics:

(1) Reliable detection of critical information within a system, based on classification guides, in evaluation studies.

Insider distinguishing factor:

The description and protection of critical information is a process done "inside" an enterprise, and is therefore tailored to the unique needs of insiders of that enterprise.

Open research problems:

(1) Develop expert systems and/or rule-based approaches to the recognition of critical content
(2) Investigate statistical modeling approaches
(3) Develop means for the reliable detection of critical content
(4) Identify ground truth in recognizing critical content.

D5: Design systems for detectability. Detecting misuse by an insider is very difficult. It could be made easier if the inherent architecture and design of an information system was created with detectability of misuse in mind. This research task aims at creating system design principles that have detectability as one of their primary features.

Research objectives:

(1) Develop system architectures that channel insider misuse into enclaves, by modularizing currently wide-open distributed systems into observable and controllable enclaves. Passage among enclaves would be regulated by "gates" at which could be placed instrumentation for observation and response.

Success metrics:

(1) Ability of a system to withstand "red team" attacks by insiders against it.

Insider distinguishing factor:

The proposed design methodology is explicitly aimed at insider use. The intent is to make an insider an "outsider" to enclaves within the larger system that he or she does not have direct, immediate need to access.

Open research problems:

(1) Design of gateways internal to a system (but perhaps within a single host on that system) that partition it into enclaves that are separately controllable, but accessible with appropriate permissions
(2) Resolution of the tension between system and data redundancy (for robustness) and the concentration of critical assets within specific enclaves
(3) Strategic deployment of sensors or "tripwires" within a system whose design is based on separate (but perhaps nested) enclaves.

D6: Determine unauthorized changes due to physical access. One distinguishing characteristic of insiders is that they often have physical access (at varying levels) to system facilities. As such, they can make physical changes to the system -- such as installing temporarily another client or host computer on the network, rerouting a cable, or plugging in a modem -- that are precluded to outsiders. It is therefore important to monitor for and detect such physical changes to a system, preferably in real-time.

Research objectives:

(1) Investigate and mitigate the risks of physical access afforded to insiders
(2) Map physical network changes dynamically
(3) Audit physical changes in order to detect unauthorized changes
(4) Determine unauthorized changes to hardware, software, and the network configuration in real-time.

Success metrics:

(1) Ability to detect and interpret "red team" physical changes to a system by insiders against it.

Insider distinguishing factor:

Insiders are unique in having physical access to many aspects of a system. This recommendation specifically addresses this special attribute of insiders.

Open research problems:

(1) Develop effective, automated techniques for network mapping
(2) Real-time dynamic change detection
(3) Automatic recognition and notification of changes
(4) System profiling and modeling that can handle the dynamic conditions of systems
(5) Scalability of any proposed solution to networks comprising thousands or tens of thousands of nodes and links.

Other detection recommendations. The workshop breakout sessions focusing on detection listed several other potential recommendations, but did not have time to discuss or elaborate on them sufficiently. We list them here as areas of potential investigation:

- *Develop standard reporting mechanisms for detection engines.* If many of the recommendations in the IPT report and this document are implemented, there will be a number of sensors, detection mechanisms, and other devices reporting potential misuse of various system assets. We need a standard reporting format, or mechanism, by which these reports can be forwarded to various correlation centers. Some uniformity in reporting is vital so that the correlation and fusion functions need not deal with a growing number of disparate formats and reports.

- *Determine the cumulative effect of negligent events.* Individual incidents of misuse of information system assets may not be critical in themselves, but their cumulative effect might be. How can we measure and determine such cumulative effects? Negligence or accidental misuse may lead to insider misuse.

- *Detect outbound copies of code and intellectual property.* Perhaps the most obvious yet damaging activity an insider can engage in is simply downloading (e.g., onto a floppy or Zip™ disk) or sending (e.g., as an e-mail attachment) a file of highly sensitive code, data, or graphical or textual information. How can these acts be detected and stopped, especially when they may occur within very short periods of time – e.g., seconds?

- *Build prosecution requirements into systems.* One of the effective means of countering the malicious insider is to quickly and effectively prosecute those malicious insiders that are caught. To do so, systems must contain effective measures for collecting and protecting information that can be used in a courtroom, including means to convince a jury that the data is valid and has not been tampered with. These are stringent requirements that should be built into the very design of information systems and networks.

The detection breakout group concluded its deliberations with an "epilogue" of considerations that cut across the various specific research recommendations:

- There are social, societal, moral, and privacy implications for the recommended research directions. Attention should be paid to these considerations before operational implementation of systems resulting from the proposed research.

- To aid in the transition from research results to tested, practical solutions, consider use of the Army Research Lab's internet testbed or the DARPA Technology Integration Center (TIC) for demonstration, testing, and evaluation.

- There is a need for aggregation, correlation, and fusion of detection data resulting from implementation of our recommendations. Such fusion can create information at higher levels of abstraction. This correlation may take place at a number of levels. There is also the need for standards and schema for reporting detection data so that such correlation can take place effectively. (DARPA is currently funding several efforts. This topic is particularly well-suited to a future workshop.)

Response

A breakout session devoted to R&D on response measures was conducted on the third day of the workshop. Among the response activities it considered were:

- *Recording, or logging:* The need to correlate physical and electronic alarms, to correlated multiple event logs (such as telephone logs, attendance records, ...), and to tailor audit logs so that they focus on the insider threat.
- *Render a damage assessment:* determine what the reach of the insider misuse has been, strictly from a system-technology viewpoint (versus counter-intelligence/espionage)
- *Reporting:* the requirement to report incidents to a central collection facility (i.e. - ARCERT, AFCERT, FIWC, DISA ASSIST, FEDCIRT) for fusion and enterprise scale assessment.
- *Repair (immediate actions):* The ability to investigate an incident from offsite or offline; the ability to intensify auditing, and do it "out of band;" the capability for autonomic repair, although this may be problematic if the insider understands how you may respond. An auto response (since it can be abused by a malicious insider) should be reserved to preserve the objective functionality of extremely critical systems only. Another potential repair action is the institution of deception and diversion measures, such as honeypots, in order to minimize damage and learn more about the intentions and level of sophistication of the malicious user. It was noted that if autonomic response is used, it may harm audit trails and records that allow the tracking, identification, and conviction of a perpetrator; therefore the actions taken by the autonomic response should be judged to be more important than catching and convicting the perpetrator.
- *Restoration* of critical services and functionality as rapidly as possible. (This has vulnerability implications for the knowledgeable insider. Restoration may create a denial of service event in order to exploit vulnerabilities in the restoration process for which the insider is given unique opportunities.)
- *Remedial systems security engineering* to more permanently repair vulnerabilities and/or remove opportunities presented by flawed process constructions.

It was noted that, at present, there are a large number of DoD and related organizations involved with various aspects of response to an event. A DoD vugraph chart used during this discussion is shown in Figure 2, listing 34 different Defensive Information Operations organizations in eight categories.[4]

DIO Organizations and Activities Study

35 Organizations Assessed

Protection	CERTs	Network Operations	Support
• Joint Task Force - Computer Network Defense • US Space Command • National Infrastructure Protection Center	• Air Force Computer Emergency Response Team • Army Computer Emergency Response Team • Navy Computer Incident Response Team • Defense Logistics Agency CERT • National Security Agency (X Group) • Carnegie Mellon University CERT/CC	• Air Force Network Operations Center • Army Network Systems Operations Center • Naval Computer and Telecommunications Command • Global Network Operations Security Center	• Joint Command and Control Warfare Center • Joint Spectrum Center • DoD Computer Forensics Laboratory • Defense Advanced Research Projects Agency • Joint C4ISR Battle Center • Army Research Lab
IW	**LE/CI**	**Intelligence**	**Other**
• Air Force Information Warfare Center • Land Information Warfare Activity • Naval Information Warfare Activity • Fleet Information Warfare Center • Information Operations Technology Center	• Air Force Office of Special Investigations • US Army Criminal Investigation Directorate • US Army Military Intelligence • Naval Criminal Investigation Service • Defense Criminal Investigative Service	• Joint Staff - J2 • Defense Intelligence Agency • Air Intelligence Agency	• National Aeronautics and Space Administration • Joint Warfare Analysis Center

Figure 2: Defensive Information Operations Organizations and Activities
Source: U.S. Department of Defense

It was noted by the response group that a responsive investigation process may either be *covert*, or *overt*. Its scope may be *enterprise-wide*, or *local.* How those two attributes are used in a particular response affect the nature of the response, and may affect the type of research needed to achieve all possible combinations in those 2x2 resultant possibilities.

- Specific response R&D recommendations are the following.

[4] The chart says "35 organizations...", but 34 are listed.

R1: Develop a capability for monitoring privacy-enhanced systems, such as those using encryption. Information systems will increasingly use internal techniques for security and privacy, such as file and link encryption. Such facilities restrict the ability to monitor systems for misuse, especially misuse by insiders capable of accessing those privacy-enhancing techniques. Research is needed on means to monitor system activities effectively when such techniques are in use. See "*Detect outbound copies of code and intellectual property"* in the section on Detection above.

Research objectives:

(1) Give analysts and investigators the ability to inspect encrypted information content during an insider incident.

Success metrics:

(1) Ability to monitor throughout all known encryption infrastructures.

Insider distinguishing factor:

Insider use of overtly-covert techniques (encryption) disables auditing of potentially unauthorized information flows from a victim organization.

Open research problems:

(1) Develop universal decryption tools to aid in forensic analysis of insider misuse incidents.

R2: Incorporate practical autonomic[5] system response into production systems. As mentioned above, insiders can carry out destructive acts in seconds. Spontaneous system responses to detected incidents may at times be required in order to thwart or mitigate damage in similar time scales – which would imply use of automated response procedures. The question is: How can these be developed so that they can be effective while not causing more harm than good. In particular, if a malevolent insider knows such an autonomic system is in use, it could be used to harm the system itself (e.g., by triggering actions that impede or disable portions of system performance). If system developers built capabilities relevant to response into system architectures, it would be much easier to initiate effective response, perhaps autonomically, within critical information systems. What are the capabilities that are relevant? Are they hard or

[5] Autonomic: "Due to internal causes or influences; spontaneous." *Webster's New Collegiate Dictionary*. New York: G.C. Merriam Co. 1980.

easy to implement? What can be done to provide incentives for their introduction into the commercial off-the-shelf systems on which DoD depends? See *"Build prosecution requirements into systems"* in the section on Detection above.

Research objectives:

(1) Create environmentally aware management technology that can dynamically modify privilege authorizations and exposure to risk.
(2) Ensure that such technology cannot be spoofed by a knowledgeable insider (other than security ad or sys ad role)
(3) Develop inherent insider threat response mechanisms in the functional applications that compose the defense information infrastructure. These inherent mechanisms themselves need to be resistant to potential insider misuse.
(4) Improve the general survivability of software products, and the software development process (developers and patch-builders as a form of insider threat).

Success metrics:

(1) Demonstrate ability to dynamically control access to all system resources (data, applications, network connectivity etc) and operating system mediated functionality.
(2) Demonstrate ability to tune technology to achieve high-fidelity sensitivity to anomalous insider activity.
(3) Develop a specification, mandated in policy, for functional requirements that aid the prosecution of potential insider misuse of the design application.
(4) Require developer, operator/administrator and user credentials to build and access DoD systems; these credentials should be encoded in system configuration and control information (e.g. - "this application was patched, administered, debugged etc. by number 8")

Insider distinguishing factor:

Insider threat creates distinguished signatures/patterns of misuse, against raised threshold parameters with respect to an interpretation of anomalous (out of character/profile) behavior. Insider threat may include system facilitators (operators, developers, admin, security personnel etc) who regularly access victim applications or information environments. A method to track the incidence of "touch".

Open research problems:

(1) Identify insider misuse characteristics with respect to various operating environments, topologies, applications and roles.
(2) Automatic recognition and notification of changes
(3) System profiling and modeling that can handle the dynamic conditions of systems
(4) Watermark and digital signature technologies to tag artifacts as evidence in insider misuse investigations.

R3: Develop data correlation tools, including data reduction for forensics, and visualization tools focused on internal misuse. In responding to a potential or actual misuse incident, data from multiple sources must be correlated and analyzed very quickly, data captured and stored for forensic analysis, and complex patterns visualized and analyzed to obtain understanding so that response can proceed. Such tools are currently lacking, especially when confronted with an insider threat – in which case many traditional sources, such as those associated with a firewall, may be unavailable or irrelevant.

Research objectives:

(1) Create multi-medium repositories that can store and index data related to insider misuse characteristics, specific misuse incidents, personnel records, telephone logs etc.

Success metrics:

(1) Formally validated insider misuse model (characterization schema) attained by mapping theoretical set of misuse activities to specific data elements.

Insider distinguishing factor:

Insider incidents often provide real-time opportunity for law enforcement agencies to discover the identity of the insider acting maliciously. Apprehension and successful prosecution (truly quantifiable threat reduction) of insiders who perform unauthorized activities requires the rapid accumulation and analysis of locally available data from all sources.

Open research problems:

(1) Develop insider misuse characterization schema that encompasses all relevant aspects of the DoD information environment.
(2) Create information systems that correlate and fuse various data sets related to insider phenomena and threat to system survivability.
(3) Demonstrate capability to correlate event-specific information.

R4: Develop a capability for surveillance of non-networked components. Not all information "system" components are on-line and capable of being queried for relevant data. Telephone logs, timecard data, various hardware settings, and so on may be relevant to determining a response, but be unavailable to correlation and fusion activities. To enable rapid and relevant response, it is highly desirable that all relevant data be collectable in real-time.

Research objectives:

(1) Incorporate multi-dimensional analysis capability in insider misuse oriented information assurance technology.

Success metrics:

(1) Demonstrate ability to dynamically access offline information for online, distributed analysis of insider misuse incident.

Insider distinguishing factor:

Insider footprint spans several technology mediums that are not normally accessible in local investigative processes.

Open research problems:

(1) Analyze the insider footprint and map sources of insider misuse evidence to the characterization schema recommended to be developed above.

R5: Consider deception technologies specifically applicable to the insider threat.

Deception has long played a vital role in military operations -- both offensive and defensive. It has tantalizing capabilities to aid in determining aspects of insider misuse that are otherwise hard to discover – in particular, the interests and intentions of a malicious insider, and that person's level of sophistication in understanding and using system facilities. Research should be conducted on the appropriate uses of deception to aid in understanding attributes of the malicious insider, and to divert him or her from critical system files and processes.

Research objectives:

(1) Develop deception techniques for information systems tailored to discovering malicious activities by insiders, and determining their intentions and level of sophistication
(2) Develop policies and procedures guiding the appropriate use of deceptive techniques in information systems.

Success metrics:

(1) Successfully discover malicious insider activity
(2) Determine intentions and level of sophistication of malicious insiders
(3) Successfully divert insiders from inappropriate access to critical files and processes.

Insider distinguishing factor:

Use of deception in information systems is one of the few ways known for discovering malicious insider activities, and determining interests and intent of those activities.

Open research problems:

(1) What system aspects (e.g., files, processes, network connections) are amenable to the introduction of deceptive techniques?
(2) How can such techniques be introduced into an operational system without having a substantial negative impact on its use?

(3) Can these techniques be used to discover misuse by highly trusted individuals, such as sysadmins, without their knowledge? How large a community must be entrusted with the knowledge that deceptive techniques are in use?
(4) Can such techniques be installed and used in such a manner that they cannot be "turned" and used against the system facilities and users they are meant to protect?
(5) What are the legal implications of using deceptive techniques in information systems?

Bibliography

NTISSIC draft, *Advisory Memorandum on the Insider Threat to U.S. Government Information Systems (IS)*, in pdf and Word formats. This was deemed essential reading for participants before the workshop.

DoD Insider Threat Mitigation Plan: Final Report of the Insider Threat Integrated Process Team, June 1999 FOUO. Essential reading before the workshop.

NIST bulletin, *Threats to Computer Systems*, March 1994

Neumann, Peter. *The Challenges of Insider Misuse*. August 1999

Attachment 1: Synopsis of Comments to Workshop by Jeffrey Hunker, National Security Council

Notes by Richard Arnold and Sara Matzner regarding
Jeffrey Hunker's dinner address to the workshop, August 17, 1999

[Mr. Hunker used three overhead transparencies as visual aids. The first provided contact information. The second was titled "Fundamental Changes" and is reproduced as item 1. below. The third was titled "Challenges" and is reproduced as item 6. below. These notes include paraphrases and direct quotes.]

1. Fundamental Changes

National Security * New Definition
Defining Threat & Response
New Organizational Models
New Roles and Missions

2. National Security * New Definition

2.1 Why are We Concerned About Cyber Attacks?

"We are dependent on Information Technology (IT) as a nation."
1/3 of US economic growth since 1995 springs from IT.
"Most operating systems that are the basis for our IT economy have vulnerabilities."
"We know that there exist hostile nations developing Cyber attacks against the US."
Russia and China have announced in the public domain
Terrorists
Organized crime
Transnational organizations
Even government Y2K software fixes were outsourced to foreign countries, and not checked when reentered into our systems. This presents an immediate and essentially impossible-to-counteract threat, which may become apparent on the new year.

2.2 "We need a New View of National Defense and Security."

Future national security threats are asymmetrical

Cyber, Chemical, Biological, and other future weapons

We've all heard of the new book from China

To compete with the US, they can't fight on our terms

They will use asymmetrical weapons

Asymmetrical weapons cost less than traditional weapons

Other constituencies will be attacked

Not just military forces, or national security targets

US Homeland and Industry are now the targets

2.3 "What are US Foreign Policy Objectives in a world of offensive Cyber capabilities?"

"This question is still a difficult challenge."

3. Defining Threat & Response

3.1 How do we define the threat and our appropriate response?

What's an Attacker?

What motivates the Bad Guys?

3.1.1 What's an "attacker"?

There is a broad spectrum of potential attacker activity

a. The "Electronic Pearl Harbor"

b. Use of Cyber tools to evade economic sanctions

c. Executive Action, to pressure a selected individual

Cyber attacks on your records … you awake to find you didn't go to the college you thought you did, etc.

d. Use of Cyber tools to "take down" 911 service in parallel with a terrorist act

"We don't have a full understanding of our potential adversary's doctrine."

3.1.2 What motivates the Bad Guys?

Cyber attack provides a unique opportunity to combine financial gain with political goals

E.g., Capture and use of CPI data prior to public announcement could result in both (1) significant monetary gain to the attacker and (2) rapid drop in confidence in the US financial system

Internationalization

Small countries can achieve status with minimal investment

Bulgaria has an internationally recognized "Lab of Virology"

US is sending software to be fixed for Y2K overseas, e.g., India

We have little Quality Control of what comes back!

"I (Hunker) strongly support investment of money in AI tools to scan for malicious code."

3.2 We need Institutional Capabilities to handle this threat

3.2.1 Intelligence

"The existing infrastructure for intelligence is not well suited to dealing with small groups of people doing things we don't understand."
We do not have systems for advanced warning of the threat.
The FBI is starting … NIPC … Information Sharing and Analysis Centers
Need: Intelligence Analysis tools!

3.2.2 US Offensive Capabilities - "if such capabilities exist … "

"It would be helpful to have insight into offensive capabilities."
"If they exist, there exists a good firewall … "
With which we could protect ourselves
"But there is not an existing policy decision for this. Is this the best way?"

3.2.3 Response and Reconstitution

"Since the reduction in the nuclear threat, there does not exist a nationwide capability to respond to a nationwide attack …
Other than FEMA, at a regional level."
"There does not exist, in the Federal Government, a Cyber response to deal with the problem; nor does there exist a national system for reconstitution, or triage if that becomes necessary."
"FEMA does not have, and has made a policy decision to not want, a cyber-capability."

3.3 Response

"It is hard to distinguish where attacks are coming from and why they are taking place."
Distinction between legal and national security tracks

"These fundamental changes will also apply to future unknown threats."

4. New Organizational Models

4.1 US Government

Speed
Technology proceeds on Internet Time; US budget works on the Calendar Year
"Our understanding of the Threat is much better than it was 6 months ago."
Diversity of Stakeholders
"This cuts across 17 agencies, and 8 to 20 major economic sectors."
Dick Clarke is a National Coordinator, not a "Cyber Czar"
[Comparison with Drug Czar authority]
"This also cuts across Congress … 13 appropriations committees are involved."
"A 14th committee was established for Y2K … I hope that it will be adopted for Cyber defense."

There is a need for a new organizational model for budgeting in government. The government budget currently is based on fiscal years each one calendar year long but the Internet evolves so quickly that we get seven years of internet change per calendar year. There is a mismatch between budget and federal decision process that requires attention at the very highest level to allow more agile response.

4.2 Industry

ISAC

We are "trying hard to encourage, on a non-regulatory basis, the sharing of information."
"There are legal questions: antitrust, liability; FOIA may apply."
"We are creating a system, for the first time, to share information about a national security threat."
"We need a system, one that handles false data, that's reliable - doesn't break under attack."

5. New Roles & Missions

5.1 DoD vs. Legal

Need to define who has cognizance, under what circumstances.
When does a cyber attack become a national security threat?

National Security vs. Federal law enforcement response: "The existing infrastructure for intelligence is not well suited to dealing with small groups of people doing things we don't understand." We do not have systems for advanced warning of the threat. However the FBI is starting Information Sharing and Analysis Centers There is a definite need for Intelligence Analysis tools!

5.2 Training and Education

We need to train and educate people to provide defense against cyber attacks
There is a shortage of IT and Information Systems Security personnel
Workers, educators, R&D teams
We are advocating a "Scholarship for Service" program
USG funds 2 years of undergraduate, or graduate, IT education in return for a commitment to provide services

In the US a high percentage of faculty and student bodies in the areas of technology are foreign citizens. There is a shortage of IT and Information Systems Security personnel, workers, educators, R&D teams. We need to train and educate people to provide defense against cyber attacks. We need to create scholarships for service program to alleviate this shortage. We are advocating a "Scholarship for Service" program which would fund two years of undergraduate, or graduate, IT education in return for a commitment to provide services

5.3 Standards

Security

We need to support the development of standards for Information Systems Security.

There is a need for due care for protecting information systems to audit against and be able to sue against. We need a legally accepted definition of due diligence for protecting information systems.

"The most effective way will probably be to work with the major accounting firms."

Most auditing is now a commodity. This risk-management add-on is potentially lucrative new business for auditing firms.

Software Development

"I'm excited about the new institutional models … "

Open Source software, Linux

Companies, e.g. HP, as brokers

6. Challenges

Privacy / Civil liberties

Encryption

Institutionalization

Metrics

6.1 Challenges [Expanded discussion]

Privacy / Civil liberties

"Defending and strengthening cyber security strengthens privacy."

Encryption

"Encryption plays only a 5% role."

Institutionalization

This is an election year. There is now bipartisan support for IP issues. Hunker wants to establish the organization now that will support IP issues while there is some bipartisanship and so that the structure is in place despite coming administration changes.

Metrics

"How do we measure success?"

7. Closing Remarks and Questions

"I am very aggressively looking for money to do malicious code detection."
"OSTP R&D of $500M is pretty safe."
Steve Skolochenko (Treasury):

Q: OMB8130 requires risk management decisions; we need quantitative threat information to implement.

A: Our office should be able to provide you with what you need.

"The NIPC is charged to identify the domestic threat. We have told industry that if you create ISAC, we will provide information to you."

Attachment 2: Attendees

Adams, Robert
Air Force Information Warfare Center
250 Hall Rd #139
San Antonio, TX 78243

Alvarez, Jorge
Space and Naval Warfare Systems Center
53560 Hull Street
San Diego, CA 92152

Anderson, Robert
RAND Corporation
P.O. Box 2138
Santa Monica, CA 90407

Anderson, Karl
NSA R2
9800 Savage Road
Ft. Meade, MD 20755

Arnold, Richard
GTE GSC
1000 Wilson Blvd. Ste 810
Arlington, VA 22209

Barnes, Anthony
Army Research Lab
C41 Systems Branch, AMSRL-SL-EI
Ft. Monmouth, NJ 07703-5602

Bencivenga, Angelo
Army Research Lab
2800 Powder Mill Road
Adelphi, MD 20783

Bozek, Thomas
Office of the Secretary of Defense / C3I
6000 Defense, Rm 3E194
Pentagon

Brackney, Richard
NSA R2, R&E Bldg
9800 Savage Road
Ft. Meade, MD 20755

Christy, James
ASDC3I/DIAP
Ste. 1101, 1215 Jefferson Davis Highway,
Arlington, Va 22202

Cowan, Crispin
Oregon Graduate Institute
P.O. Box 91000
Portland, OR 97291

Dunn, Timothy
Army Research Lab
2800 Powder Mill Road
Adelphi, MD 20783

Dunphy, Brian
Defense Information Systems Agency
701 S.Courthouse Rd D333
Arlington VA

Ghosh, Anup K.
Reliable Software Technologies
21351 Ridgetop Circle, Ste 400
Dulles, VA 20166

Gligor, Virgil
University of Maryland
Electrical/Computer Engineering, AVW 1333,
College Park, MD 20742

Gilliom, Laura
Sandia National Labs
P. O. Box 5800-0455
Albuquerque NM

Goldring, Tom
NSA R23
9800 Savage Road
Ft. Meade, MD 20755

Hotes, Scott
NSA R225 R&E Bldg
9800 Savage Road
Ft. Meade, MD 20755

Hunker, Jeffrey
National Security Council
White House #303
Washington DC 20504

Jaeger, Jim
Lucent Technologies
Box 186, Columbia, MD 21045

Longstaff, Thomas
CERT/CC
4500 Fifth Avenue
Pittsburgh, PA 15213

Lunt, Teresa
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304

Matzner, Sara
U. Texas at Austin Applied Research Labs
Information Systems Laboratory, P.O. Box 8029,
Austin Texas 78713

Maxion, Roy
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

McGovern, Owen
DISA
Letterkenny Army Depot
Chambersburg, PA 17201-4122

Merritt, Larry D.
NSA
9800 Savage Road
Ft. George G. Meade, MD 20755

Neumann, Peter G
SRI International
333 Ravenswood Ave.
Menlo Park, CA 94025

Skolochenko, Steven
Office of Information Systems Security
1500 Penn. Ave. NW, Annex, Rm. 3090,
Washington, DC 20220

Skroch, Michael
DARPA/ISO
3701 N. Fairfax Dr.
Arlington, VA 22203

Solo, David
Citibank
666 Fifth Ave., 3rd Floor/Zone 6
New York, NY 10103

Teslich, Robyne
Lawrence Livermore National Laboratory
PO Box 808, Room L-52
Livermore CA 94550

Tung, Brian
USC Information Sciences Institute
4676 Admiralty Way Ste. 1001,
Marina del Rey, CA 90292

van Wyk, Kenneth
Para-Protect
5600 General Washington Drive ste. B-212
Alexandria, VA 22312

Walczak, Paul
Army Research Laboratory
2800 Powder Mill Road
Adelphi, MD 20783

Zissman, Marc
Mit Lincoln Laboratory
244 Wood Street
Lexington, MA 20420

Attachment 3: Workshop Agenda

Day 1 - Monday 16 August

1300 - 1400 Registration

1a 1400 - 1700 (Plenary) Opening Session, *Dick Brackney,* NSA, chair
Co-sponsors' opening remarks & expectations of the Workshop,
Dick Brackney, Mike Skroch, Paul Walczak
Round of introductions
1420 Initial statement, *Dick Schaeffer,* DoD videotape from the Pentagon
1430 Live interactive phone hookup with *Dick Schaeffer*
1547 Break
1607 Discussion of the Integrated Process Team report, *Tom Bozek*
1702 Review of IPT technology recommendations, *Bob Anderson*
1730 - 1930 No-Host reception at RAND

Day 2 - Tuesday 17 August

0730 - 0830 Continental breakfast

2a 0830 - 1000 (Plenary) *Peter Neumann*, SRI, Chair
0827 Summary of yesterday's meeting with John Hamre, *Angelo Bencivenga*
0837 Review of DoE's Wash DC Workshop on Insider Threats, *Dick Brackney*
1000 - 1015 Break

2b 1015 - 1100 (Plenary) Prevention, *Mike Skroch*, DARPA, Chair
What improvements in prevention are needed to enhance the detectability of insider misuse?
What are the limiting factors today, and what must change?
What are the current expectations of future R&D re: insider misuse?
1100 Scoping the workshop: seeking new technical R&D directions
Scoping the insider misuse problem
Classes of insiders and classes of insider misuse
Threats to be addressed: insiders in context; outsiders becoming insiders? accidental misuse?
How does insider misuse differ from outsider misuse?
Decomposition of the problem
Discussion and organization of effective breakout structures most suited to the diversity of the attendees

2c 1130 - 1200 (First set of breakout groups)
Insider threats and vulnerabilities, *Paul Walczak*
Prevention of insider misuse, *Mike Skroch*
Detection of insider misuse, *Dick Brackney*

2c 1300 - 1500 (First set of breakout groups, continued)
1500 - 1515 Break

2d 1515 - 1700 (Plenary) Summary of each of the day's 2c breakout groups

and further discussion
1700 Workshop concludes for the day
1800 Informal Workshop banquet at Miramar Sheraton Hotel
Banquet speaker, *Jeffrey Hunker,* National Security Council

Day 3 - Wednesday 18 August

0730 - 0830 Continental breakfast
3a 0830 - 0930 (Plenary) open discussion
0930 - 1000 Breakout discussion begin
1000 - 1015 Break
3b 1015 - 1200 (Breakout sessions) R&D directions
What are the R&D implications of the identified differences with respect to insider misuse
What new capabilities must be developed?
What are the highest-priority directions for the future, and why?
Prevention including Misuse Tolerance, *Mike Skroch, Crispin Cowan*
Detection, *Dick Brackney, Brian Tung*
Response, *Paul Walczak*
1200 - 1300 Lunch
3c 1300 - 1500 (Plenary) Summaries of 3b breakouts, *Dick Brackney*, chair
1500 - 1515 *Willis Ware* analysis of the summaries
1515 - 1530 Break
3d 1530 - 1630 (Plenary) Discussion of open issues and next steps forward
3e 1630 - 1645 (Plenary) Closing remarks
1645 Workshop concludes